

STREAMING MEDIA

FIELD OF THE INVENTION

- 5 The invention relates to streaming of media from a streaming server to a mobile client device over an air-interface.

BACKGROUND OF THE INVENTION

- 10 A packet switched streaming service PSS is currently being standardized for mobile environment by 3GPP (3rd Generation Partnership Project). An example of a communications system capable of streaming media (streaming video and/or audio) is shown in Figure 1. The system comprises a streaming server 111 which is coupled to an IP-network (Internet Protocol) 104. The IP-network 104 may be, for example, the Internet or a service provider operator's intranet (an intranet network
15 belonging to the operator's domain). The IP-network 104 is coupled to a core network 103 of a mobile communications network via a G_i interface. The mobile communications network also has a radio access network (RAN) 102 coupled to the core network 103. The radio access network 102 provides mobile communications devices 101 with access to the mobile communications network over an air-
20 interface. Said access may be provided either by circuit switched means (circuit switched voice or data call) or packet switched means or both. In the following, GPRS (General Packet Radio Service) is used as an example of a packet switched means to communicate over the air-interface. Concerning GPRS, the generic term
25 radio access network (RAN) is considered to comprise base (transceiver) stations (BS) and base station controllers (BSC).

- In streaming media a sequence of 'moving images' (i.e. video) or sound (i.e. audio) or a sequence of 'moving images' together with sound (i.e. multimedia) is
30 sent in a compressed form from the streaming server 111 to a mobile communications device 101 (hereinafter referred to as a client device 101). In contrast to the

technique in which an entire media file has to be arrived at the client before it can be played, the streaming technique enables sending of media (video and/or audio) from the streaming server 111 to the client device 101 in a continuous manner and playing of the media as it arrives at the client.

5

When media is streamed in cellular mobile communications systems such as the system presented in Figure 1, new problems specific to the mobile environment arise. These problems are mostly due to the different restrictions of mobile systems. One such problem arises in a cell reselection (i.e. handover) situation presented in Figure 2.

10

In Figure 2, a first base station BS1 of the radio access network 102 serves the area of a first cell 201 whereas a second base station BS2 of the radio access network 102 serves the area of a second cell 202. Before cell reselection, the client device 101 is served by the first base station BS1, i.e. the client device 101 has an active radio link with the first base station BS1. This situation is shown on the left side of the Figure 2. After the cell reselection, the client device 101 is served by the second base station BS2. The radio link with the first base station BS1, shown by a dotted line, is dropped. This situation is shown on the right side of the Figure 2.

15

20

The cell reselection (CR) can be divided in three periods in time:

- a) Pre-CR period
- 25 b) CR period
- c) Post-CR period.

During the pre-CR period, received signal quality in the first cell weakens and the client device 101 starts cell reselection signaling. The client device 101 is able to receive streaming media via the first base station BS1 during that period. During the CR period, the actual cell reselection is performed. The client device 101 can

30

not receive streaming media during that period. During the post-CR period, the client device 101 is able to receive streaming media via the second base station BS2.

5 The cell reselection may cause a long break in service. For example, a cell reselection (CR period) in a case where GPRS is used as radio bearer may take up to 30-40 seconds. This will affect an ongoing streaming session. For example, part of the packets carrying streaming media may be lost and/or freezing of the played media may occur at the client device 101. If the streaming media comprises a
10 video stream, the freezing means that a still image will appear and stay for a period of time on a client device display before the streaming media is again received and can be played again. If the streaming media comprises an audio stream, the freezing means that no sound is played back (i.e. silence) until the CR period is over.

15

Accordingly, there is a need to find appropriate means for reducing the effects of cell reselection on an ongoing streaming session.

SUMMARY OF THE INVENTION

20

According to a first aspect of the invention, there is provided a method for streaming of media from a streaming server to a mobile client device over an air-interface, wherein the method comprises:

requesting the streaming server to send streaming media which the mobile client
25 device is not able to receive due to a cell reselection.

30

A request requesting the streaming server to send the streaming media may be sent from the mobile client device to the streaming server before or after the mentioned cell reselection.

The mentioned request is not limited to requesting to send only the media which

the mobile client device is (or was) not able to receive due to the cell reselection (during a cell reselection period). In an embodiment, the request may be construed such that what is actually requested is: resending of a portion of a media content not received as well as sending any remaining portion of the media content.

5

The term media is considered to mean either video or audio or another media, such as still image, or any combination thereof, i.e multimedia.

Preferably, a resending request is sent in response to the cell reselection. The re-
10 sending request is preferably an application layer request sent between a mobile client device application and a streaming server application.

In a preferred embodiment, the streaming server is requested to send at an in-
creased speed for a determined period of time in order to increase at the client
15 device a degree of fullness of a temporary store, such as a buffer, in which the streaming media is temporarily stored before playback. Preferably, this is done by by requesting the streaming server to switch from sending a higher bit rate media stream to sending a lower bit rate media stream at an increased speed. The higher and lower bit rate streams are preferably provided by a multi-rate codec.

20

According to a second aspect of the invention, there is provided a mobile client device for receiving streaming media from a streaming server over an air-
interface, the mobile client device comprising:

means for requesting the streaming server to send streaming media which the mo-
25 bile client device is not able to receive due to a cell reselection.

Preferably, the air-interface couples said mobile client device to a mobile commu-
nications network. The mobile client device preferably comprises a cellular mo-
bile phone.

30

According to a third aspect of the invention, there is provided a streaming server

for sending streaming media to a mobile client device over an air-interface, the streaming server comprising:

means for receiving a request requesting the streaming server to send streaming media which the mobile client device is not able to receive due to a cell reselection; and

5 tion; and

means for acting upon the received request.

According to a fourth aspect of the invention, there is provided a system comprising a streaming server and a mobile client device, for streaming of media from the streaming server to the mobile client device over an air-interface, the system comprising, at the mobile client device:

10

means for requesting the streaming server to send streaming media which the mobile client device is not able to receive due to a cell reselection, the system further comprising, at the streaming server:

15 means for receiving the request; and

means for acting upon the received request.

According to a fifth aspect of the invention, there is provided a computer program executable in a mobile client device, the computer program comprising:

20 program code for causing the mobile client device to request the streaming server to send streaming media which the mobile client device is not able to receive due to a cell reselection.

According to a sixth aspect of the invention, there is provided a computer program executable in a streaming server, the computer program comprising:

25

program code for causing the streaming server to receive a request requesting the streaming server to send streaming media which the mobile client device is not able to receive due to a cell reselection; and
program code for acting upon the received request.

30

Dependent claims contain preferable embodiments of the invention. The subject

matter contained in dependent claims relating to a particular aspect of the invention is also applicable to other aspects of the invention.

BRIEF DESCRIPTION OF THE DRAWINGS

5

Embodiments of the invention will now be described by way of example with reference to the accompanying drawings in which:

- Figure 1 shows a communications system capable of streaming media;
- 10 Figure 2 shows a cell reselection situation in the communications system of Figure 1;
- Figure 3 shows a client device buffer in a situation in which the client device buffer is shorter than the time that cell reselection takes;
- 15 Figure 4 shows a client device buffer in a situation in which the client device buffer is longer than the time that cell reselection takes;
- 20 Figure 5 shows another illustration of the situation which is behind Figure 4;
- Figure 6 shows preferred action in a preferred embodiment of the invention;
- Figure 7 illustrates a preferred embodiment of the invention;
- 25 Figure 8 illustrates a preferred embodiment of the invention;
- Figure 9 shows a client device according to a preferred embodiment of the invention;
- 30 Figure 10 shows a streaming server according to a preferred embodiment of the

invention; and

Figure 11 illustrates three ways of sending messages according to preferred embodiments of the invention.

5

DETAILED DESCRIPTION

The system shown in Figure 1 can also be used in the preferred embodiments of the invention (the same applies to Figure 2). Accordingly, the system comprises a
10 streaming server 111 which is coupled to an IP-network 104. The IP-network 104 may be, for example, the Internet or a service provider operator's intranet. The IP-network 104 is coupled to a core network 103 of a mobile communications network via a G_i interface. The mobile communications network also has a radio access network (RAN) 102 coupled to the core network 103. The radio access net-
15 work 102 provides mobile communications devices 101 with access to the mobile communications network over an air-interface. Said access may be provided either by circuit switched means (circuit switched voice or data call) or packet switched means or both. In the following, GPRS (General Packet Radio Service) is used as an example of a packet switched means to communicate over the air-interface.

20

In order to better understand the preferred embodiments of the invention, let us consider the following example.

In this example, a mobile communications device 101 (again referred to as a client
25 device 101) has established a streaming session with the streaming server 111 when being served by the first base station BS1. In this example and in the following description, mostly a video streaming session is considered. However, a corresponding consideration applies to an audio streaming session as well as to a multimedia (e.g. video with audio) streaming session.

30

RTSP (Real Time Streaming Protocol) protocol is used in the streaming session

setup. Once the session has been established, the streaming itself may be performed (i.e. media flow may be sent) according to RTP (Real time Transport Protocol) or another protocol. However, if it is desired to make a change in the established session, this will again be done by using RTSP.

5

At the beginning of the established streaming session, the received streaming media (media flow), transmitted in packets (or frames), is buffered (i.e. temporarily stored) in a buffer at the client device (hereinafter referred to as a client device buffer). When the client device buffer is full, playing (playback) of the streaming
10 media is started at one end of the buffer while the buffer is continuously filled at the other end with received streaming media. Therefore, the buffer should remain nearly full during normal operation of the system.

The client device now undergoes a cell reselection from the first cell 201 to the
15 second cell 202 (Fig. 2). It is to be noted that, although the radio access network (RAN) 102 is illustrated in Figure 2 by a single cloud only, it may comprise different radio access networks. It may comprise a GPRS (GPRS-only) RAN, a 3G (third generation) RAN or a combination thereof. The first cell 201 (served by the base station BS1) and the second cell 202 (served by the base station BS2) may
20 therefore be, for example, cells of the GPRS RAN or 3G RAN. Accordingly, the cell reselection (handover) can occur, for example, between base stations belonging to GPRS RAN, between base stations belonging to 3G RAN or between a base station belonging to GPRS RAN and a base station belonging to 3G RAN.

25 The client device 101 can not receive streaming media during the CR period. Therefore, when the CR period begins the client device buffer starts emptying because the buffer content is further played while the continuous filling is stopped. Depending on the size of the client device buffer, two different cases can be identified:

30

i) The client device buffer (in time) is shorter than the time that cell reselection

takes, i.e. the time to empty a full client device buffer is shorter than the CR period,

- ii) The client device buffer (in time) is longer than the time that cell reselection takes, i.e. the time to empty a full client device buffer is longer than the CR period.

In the first case (Figure 3), the buffer will become totally empty during the CR period. Since the buffer becomes empty, the client device 101 has no media to play anymore. Therefore, the client device 101 starts rebuffering. Rebuffering can be started during the post-CR period. In order to get the rebuffering started the client device 101 can use an RTSP PAUSE/PLAY method. In this method, the client device 101 first sends an RTSP PAUSE message to the streaming server 111. The PAUSE message will cause the server 111 to pause sending of the media flow. Subsequently, the client device 101 sends an RTSP PLAY message which requests the streaming server 111 to start resending the streaming media from the last received frame onwards. Rebuffering starts when the client again starts receiving the streaming media. When the client device buffer is full, playing of the streaming media is restarted from the last received frame onwards. In the meantime, the last received frame is shown as a still image on a client device display. However, in spite of the drawback of showing a still image, no packets are effectively lost since the media flow which is not received at the client device during the CR period is resent from the server 111 to the client device 101 after the CR period.

25

In the second case (Figure 4), the whole content of the buffer has not been played during the CR period. Therefore, the buffer only becomes partially empty during the CR period, i.e. there is still streaming media (DATA) in the buffer when the CR period ends. During the post-CR period (after the CR period) playing of the streaming media stored in the buffer before the CR period (i.e. DATA*) is further continued at one end of the buffer while the other end is filled with the streaming

30

media received after the CR period (i.e. DATA^{**}). If nothing is done, DATA^{*} will eventually end and playing is continued with the first frame of data DATA^{**}. No still image is shown. However, the first frame of DATA^{**} is not the frame that should be played right after the last frame of DATA^{*}. Since the client device 101
 5 could not receive streaming media during the CR period it has not even received the frame that should be showed after the last frame of DATA^{*}. There exists a temporal gap in the buffered media. As a result, the client device 101 will skip showing the sequence of streaming media (video) which could not be received due to the cell reselection. Concerning an audio stream, the same effect occurs, i.e.
 10 the client device 101 will skip playing sound which could not be received.

The second case is further illustrated with the aid of Figure 5. It is to be noted that Figure 5 only shows a very simplified case. The time is considered to flow from left to right in Figure 5. The streaming media sequence that the server sends comprises video frames A-S. The frames are received at the client device and buffered
 15 before playing. The playing of the media starting from frame A starts when the buffer is full (i.e. when the frames A-G have been stored in the buffer), the frames (moving video picture) are subsequently shown on a client device display. The buffering of the frames following frame G is not shown in Figure 5 for clarity
 20 reasons.

The media flow from the server to the client device is not possible during the CR period. Therefore, the client device does not receive the frames L-O. Otherwise, the media flow is perfectly received at the client device. Now that the the
 25 CR period is shorter than the buffer size, the showing of the frame (frame K) last received before the CR period will be followed by the showing of the frame (frame P) first received after the CR period. Video frames L-O are not shown at all. Concerning an audio stream, the same effect occurs, i.e. the user hears a pause and a temporal gap in the sound.

30 There exists two type of frames: intra frames and inter frames. The intra frames

contain all necessary information of an image whereas inter-frames only contain changes or predicted changes compared to a previous image. The previous image may be an intra frame or an inter frame. Therefore, if the frame P is an intra frame the shown picture only warps ahead in time from frame K to frame P. In this case
5 the frames L-O are not simply shown to the user. But if the frame P is an inter frame severe distortion in the shown (moving) picture is likely to occur before the next intra frame is received and played.

A preferred embodiment of the invention as shown in Figure 6 concentrates on the
10 above mentioned problem relating to the second case. In this embodiment, the buffer is longer than the duration of the CR period. Therefore, the buffer does not become totally empty during the CR period but only partially empty. It is to be noted that Figure 6 only shows a very simplified case. The time is considered to flow from left to right in Figure 6. The streaming media sequence that the server
15 sends comprises frames A-S. The frames are received at the client device and buffered before playing. The playing of the media starting from frame A starts when the buffer is full, the frames (moving video picture) are subsequently shown on a client device display. The buffering of the frames following frame G is not shown in Figure 6 for clarity reasons. The media flow from the server to the client
20 device is not possible during the CR period. Therefore, the client device does not receive the frames L-O. Otherwise, the media flow is perfectly received at the client device.

When the CR period is over the client device knows exactly what is the last frame
25 that it has received before the CR period (i.e. during the pre-CR period). In this embodiment it is frame K. The client device then requests, right after the CR period has been ended, the streaming server to start resending the streaming media from the last received frame onwards. Upon receiving the request the streaming server starts resending the media flow. The first frame to be sent is in
30 this embodiment frame L. Any frames sent and received after the CR period but before the resending has started (possibly some frames starting from frame P) are

ignored by the client device. When the streaming media (of frames A-K) stored in the buffer before the CR period ends, then the playing is continued with the first frame (frame L) of the resent streaming media received after the CR period. Therefore, no discontinuity in the shown video picture should appear to the user.

- 5 In Figure 6, the resent video frames L-O which would have otherwise failed to have been received at the client device are shown in bold font. Correspondingly, concerning an audio stream, the streaming server is requested to start resending the audio stream at a point at which reception has been stopped in the client device. No discontinuity in playback should appear.

10

- In practice, the resending request can be performed by the RTSP PAUSE/PLAY method. In this method, the client device first sends, after the CR period has been ended, an RTSP PAUSE message to the streaming server. The PAUSE message will cause the server to pause sending of the media flow. However, the playing of the streaming media received at the client device is not paused unless the buffer gets empty (which should not be the case). Subsequently, the client device sends an RTSP PLAY message to the streaming server 111. The PLAY message contains information on the starting point of the resending. At the end of the CR period the client device knows the time of the last received frame. Based on that, the client device determines the starting point before sending the PLAY message. The PLAY message causes the server to start resending.
- 15
- 20

An example of the PAUSE message is as follows:

25 PAUSE rtsp://example.com/foo RTSP/1.0
CSeq: 6
Session: 354832

- The PAUSE message informs the server that a change is coming. An example of the subsequently sent PLAY message is as follows:
- 30

PLAY rtsp://example.com/foo RTSP/1.0

CSeq: 7

Session: 354832

Range: npt=28.00-

5

The message field 'Range' tells the starting point of resending. In this example, the starting point is 28 seconds from the start of the streaming media sequence. In relation to the embodiment shown in Figure 6, this point would be exactly the time of frame L in the streaming media sequence.

10

In the embodiment just described, the frames (packets) that were lost during the CR period will be resent from the server to the client device. Also, since the buffer size is longer than the duration of the CR period, no interruption in the shown video picture should appear and the user experience is maximized.

15

However, as the case is such that during normal streaming the degree of fullness of the buffer does not change, since the buffer is emptied (played) at the same rate as it is filled, this may cause an additional problem. In the embodiment just described, after the cell reselection, the client device buffer stays emptier than before, for the reasons as described in the foregoing. Therefore if, for example, a new cell reselection is to be performed in a near future, the emptier buffer may cause the same drawback as discussed in connection with the first case above (i.e. the case in which the client device buffer was shorter than the time that cell reselection takes).

25

A preferred embodiment of the invention concentrates on the above identified problem. In this embodiment, to ensure smooth behaviour of the shown video picture (correspondingly: played sound), the client device buffer is filled after the CR period, for a period of time, at a higher rate than it is emptied (played). This period may be called a filling period. When the filling period is over the buffer is full again and normal streaming in which the buffer is filled at the same rate as

30

emptied is resumed.

It is to be noted that normally raising the degree of fullness of the buffer would require the playing to be paused. In the present embodiment, playing is not paused
5 but the degree of fullness of the buffer can still be raised due to the clever buffer management in which the buffer is filled, during playing, at a rate higher than the playing rate.

To raise the degree of fullness of the buffer without pausing the playback, the client device, in this embodiment, requests the server to switch to sending a lower bit
10 rate streaming media sequence, but to use in the actual sending the same transmission bit rate (hereinafter referred to as an original transmission bit rate) as before. In order to reach the original transmission bit rate the client device requests the server to speed up the transmission of the lower bit rate sequence by a speeding
15 factor. Speeding up the transmission bit rate causes more data to be written in the buffer than read out from the buffer. Therefore, the degree of fullness of the buffer increases as desired.

In other words, the server is asked to switch from sending the original sequence
20 encoded at a first bit rate to sending a new corresponding sequence encoded at a second bit rate lower than the first bit rate and to increase the transmission bit rate of the new sequence in order to reach the original transmission bit rate (bandwidth). It is to be noted that the bit rate at which a media stream is encoded (and will be decoded) is a different concept than the transmission bit rate. The bit rate
25 at which the media stream has been encoded has an effect on the picture quality. If the media stream has been encoded at a higher bit rate, this means that more bits have been used in encoding compared to encoding at a lower bit rate. This typically results in a better picture quality. The transmission bit rate, on the other hand, is the bit rate at which the media stream is actually sent and it depends on
30 the available bandwidth.

The request to switch to sending a lower bit rate sequence and to speed up the transmission bit rate can be delivered by using the RTSP PAUSE/PLAY method. When the client device desires to fill the client device buffer during the post-CR period, it first sends a PAUSE message, corresponding to the one presented
 5 already in the preceding description, to the server to indicate that a change is coming. Subsequently it generates a PLAY message and sends it to the server.

An example of such a PLAY message is as follows:

```
10    PLAY rtsp://example.com/foo RTSP/1.0
      CSeq: 7
      Session: 354832
      Range: npt=28.00-40.00
      Bandwidth: 20000
15    Speed: 1.5
```

This message contains two optional message fields, namely 'Bandwidth' and 'Speed' to be understood at the client device and streaming server. These fields are already specified as optional fields by the IETF (Internet Engineering Task
 20 Force) in the standard RFC 2326 (Real Time Streaming Protocol).

The message field 'Bandwidth' tells the server to change into sending a lower bit rate sequence (here: a sequence whose bit rate is 20 kbps) and the message field 'Speed' tells the server to speed up the sending by a speeding factor (here: 1.5).
 25 The message field 'Range' tells the starting point and stopping point in the streaming media sequence in time units (calculated from the beginning of the streaming media sequence(s)).

The PLAY message shown would suit for an example in which the server initially
 30 sends a 30 kbps bit rate sequence at an original transmission bit rate of 30 kbps and in which the server, in order to fill the buffer, is desired to switch into sending

a sequence whose bit rate is 20 kbps and to speed up the sending at a speeding factor 1.5 to reach the original transmission bit rate of 30 kbps during a 12 seconds filling period (i.e. between the time instants of 28 and 40 seconds of the media sequence).

5

If the client knows the possible bit rate options (they are typically communicated from the server to the client during session setup, for example, with the aid of a 200 OK response to an RTSP DESCRIBE message) it can calculate the speeding factor needed to reach the original bitrate using the following formula:

$$\text{speeding factor} = \frac{\text{original transmission bit rate}}{\text{new transmission bit rate}}$$

10

When the buffer is full, the client device sends another RTSP PAUSE and PLAY message pair to the server. An example of the PLAY message is as follows:

PLAY rtsp://example.com/foo RTSP/1.0

15

CSeq: 67

Session: 354832

Range: npt=40.00-

Bandwidth: 30000

Speed: 1.0

20

This exemplary message requests the server to start sending the (original) 30 kbps bit rate media sequence (Bandwidth: 30000) at the original transmission bit rate 30 kbps (Speed: 1.0) starting from the time instant 40 seconds of the media sequence (Range: npt=40.00-).

25

The client device can calculate the length of the buffer filling period using the following formulas:

$$\text{filling period} = \frac{\text{LowSeqTime}}{\text{speeding factor}}$$

wherein

$$\text{LowSeqTime} = (\text{BufferSize} - \text{BufferData}) \cdot \frac{\text{original transmission bit rate}}{\text{original transmission bit rate} - \text{new transmission bit rate}}$$

In these formulas, *LowSeqTime* indicates the duration of the lower bit rate sequence playing time in the client device, *BufferSize* indicates the size of the buffer in seconds and *BufferData* indicates the data left in the buffer in seconds.

The filling of the buffer is illustrated in Figure 7. It is to be noted that Figure 7 only shows a very simplified case. All delays, for example, have been ignored. The time is considered to flow from left to right in Figure 7. It can be seen that the buffer starts emptying when the CR period starts. Upon receiving the first PLAY message the server switches to sending a media sequence having a lower bit rate (here: 20 kbps) and speeds up the sending by a speeding factor (here again: 1.5). Since the speeding factor is 1.5, the buffer filling takes twice the time that the emptying period lasted. For example, if the emptying period lasted for 25 seconds, the filling period should be 50 seconds with this speeding factor. When the buffer is full the second PLAY message is sent and the server switches to sending the original media sequence having a higher bit rate (here: 30 kbps) and continues sending at the original transmission bit rate.

20

In the foregoing, it has been described that, for example, a video stream typically comprises both intra and inter frames, wherein the intra frames are ‘independent frames’ containing all necessary information of an image whereas inter-frames only contain changes or predicted changes compared to a previous image. In the following preferred embodiment of the invention, the timing of the switch from the lower bit rate sequence to the original bit rate sequence is more closely de-

25

scribed with this respect.

In this embodiment the purpose is to time, whenever possible, the switch from the lower bit rate sequence to the original bit rate sequence to take place at the point of an intra frame of the original bit rate sequence (or, more generally, at the point of an intra frame of the sequence to which the sending is switched). In this way, prediction errors and/or jumps in media (frame) flow can be avoided. Figure 8 illustrates this embodiment. It is to be noted that Figure 8 only shows a simplified case. The time is considered to flow from left to right in Figure 8. In this embodiment, the first switch from the original bit rate sequence to the lower bit rate sequence is effected at point 28 seconds with the aid of a first PAUSE/PLAY message pair.

As illustrated in Figure 8, the position of intra frames in the original bit rate sequence may differ from the ones of the lower bit rate sequence. However, based on:

- the time required to totally fill the buffer, i.e. '*filling period*',
- the time distance of two adjacent intra frames in the original sequence, i.e. an intra frame rate, '*IFrameTimeOriginal*', and
- the time at which the first switch from the original bit rate sequence to the lower bit rate sequence is effected, i.e. '*SwitchTime*',

it can be calculated a suitable point in time for switching back to the original bit rate sequence. The parameter '*IFrameTimeOriginal*' can be calculated in the client device from the original bit rate sequence before switching to the lower bit rate sequence. The calculation of the suitable point in time for switching back to the original bit rate sequence can be performed using the following formula:

$$SeqChangeTime = IFrameTime_{Original} * \left\lfloor \frac{SwitchTime + filling\ period}{IFrameTime_{Original}} \right\rfloor$$

In this formula, '*SeqChangeTime*' indicates the point in time for swiching back to the original bit rate sequence and the brackets indicate a *floor*-function (or *trunc*-function) which truncates the fractional part of the value computed in the brackets.

5 For example, if the time required to totally fill the buffer, i.e. '*filling period*' is 16 seconds, the time at which the first switch is effected, i.e. '*SwitchTime*' is 28 seconds and the time difference between two adjacent intra frames in the original sequence, i.e. '*IFrameTime_{Original}*' is 5 seconds, then the point in time for swiching back to the original bit rate sequence is 40 seconds (*SeqChangeTime*=5**floor*((28+16)/5) = 40 s). Accordingly, the second
10 PAUSE/PLAY message pair is sent at point 40 seconds, as indicated in Figure 8. In this embodiment, the starting point and stopping point to be placed in the message field 'Range' of the PLAY message are 28 seconds and 40 seconds, respectively.

15

It is to be noted that, due to the intra frame rate, it might not in all cases be possible to totally fill the buffer. For example, in the example just described, it would have taken 4 seconds more to totally fill the buffer. However, the formula gives the closest suitable intra frame to the buffer filling period.

20

In another embodiment, it may well be that, at the point of the last received frame before the CR period, there is not an intra frame in the lower bit rate media sequence. In this embodiment, the starting point of the resending request is adjusted by a necessary amount of frames (or time) backwards such that there will be an
25 intra frame in the lower bit rate sequence at the starting point. In this case, a set of last received frames before the CR period, which belong to a time period subsequent to the starting point, are ignored at the client device in order to guarantee a continuous playback.

In another embodiment of the invention, the switches between two different bit rate sequences are performed without the timing method(s) just described. In this embodiment, the first switching point (from higher to lower bit rate sequence) is directly determined by the last received frame and the second switching point
5 (from lower to higher bit rate sequence) is directly determined by the the time required to totally fill the buffer (*'filling period'*). The switching points can therefore end up being at the point of an intra or inter frame. Should a switching point end up being at the point of an inter frame (for example a P-frame), a small prediction error may occur in the played media, however, in this embodiment, the
10 buffer can be totally filled.

Another way to perform switching is to use so called "switch frames". These are frames which contain "difference" information between corresponding frames in different bit rate sequences. In this embodiment, a bridge and switch between the
15 two sequences is performed with the aid of these frames using the difference information.

Figure 9 shows a client device 101 according to a preferred embodiment of the invention. The client device may be a mobile station of a cellular radio telephone network. The client device 101 comprises a processing unit MCU, a radio frequency part RF, and a user interface UI. The radio frequency part RF and the user interface UI are coupled to the processing unit MCU. The user interface UI typically comprises a display, one or more speaker and a keyboard (not shown) with the aid of which a user can use the device 101.

25

The processing unit MCU comprises a processor (not shown), a memory 210 and computer software. The software has been stored in the memory 210. The aforementioned client device buffer 240 is also comprised in the memory 210. The processor controls, in accordance with the software, the operation of the client
30 device 101, such as receiving of the streaming media sent from the server 111 and sending of requests to the server 111 via the radio frequency part RF, reading and

writing of received streaming media (video and/or audio) in the buffer 240, and presenting of the received streaming video on the display and audio on the one or more speaker of the user interface UI. A suitable size of the buffer (in time) can be, for example, 1.5 or 2 times the maximum (or average) CR period time.

5

The software comprises a streaming client software application 220 (hereinafter referred to as client software 220), a protocol stack 230 for implementing the necessary protocol layers such as an RTP layer, an RTSP layer, an SDP (Session Description Protocol) layer, a TCP layer (Transmission Control Protocol), an IP
10 layer and, below the IP-layer, lower protocol layers. In addition, the software comprises as a part of the client software 220 a media player for playing the received media.

The processor generates, based on the client software 220 the aforementioned
15 PAUSE and PLAY messages and sends them to the server 111 via the radio frequency part RF. The processor also performs, based on the client software 220, the necessary calculations relating to the speeding factor, relating to the buffer filling period and relating to the suitable point in time for switching from a lower bit rate sequence back to the original bit rate sequence.

20

The generation and sending of the resending request (PLAY message) and other appropriate action is triggered by a cell reselection event occurred in the mobile client device 101. The event can be detected by the client software 220 by reception of an asynchronous message from an API (Application Programming Inter-
25 face) provided by the lower layers of the protocol stack 230. Alternatively or in addition, the event can be detected by monitoring the buffer level, i.e. the degree of fullness of the client device buffer 240. In this case, if the buffer 240 does not receive data for a certain amount X of time (depending on the implementation, the parameter X can be defined as a constant value, and it is a threshold for the client
30 software 220 to understand that the cell reselection event has occurred), and if the client device 101 afterwards starts to receive data after a certain variable amount

Y of time (wherein $Y > X$, and Y is the real duration of the CR period), then the client can trigger the action described in this description.

5 Figure 10 shows a streaming server 111 according to a preferred embodiment of the invention. The streaming server 111 comprises a processing unit CPU, a first memory 310, an IP network interface 350, and a second memory 360. The first memory 310, the IP network interface 350, and the second memory 360 are coupled to the processing unit CPU.

10 The processing unit CPU controls, in accordance with computer software stored in the first memory 310, the operation of the streaming server 111, such as processing of requests received from the client device 101 and the sending of pre-recorded media streams, stored in the second memory (disk) 360, to the client device 101 via the IP network interface 350.

15

The software comprises a streaming server software application 320 (hereinafter referred to as server software 320), a protocol stack 330 for implementing the necessary protocol layers such as an RTP layer, an RTSP layer, an SDP layer, a TCP layer, an IP layer and lower protocol layers.

20

The PAUSE and PLAY messages sent from the client device 101 are received via the IP network interface 350. A processor (not shown) of the processing unit CPU processes the messages according to the server software 320 and the protocol stack 330 and takes appropriate action.

25

The present invention provides means for reducing the effects of cell reselection on an ongoing streaming session. It is to be noted that, according to preferable embodiments of the invention, a request for resending (e.g. a PLAY message) is sent on the application layer, i.e. between the client software application 220 and
30 the server software application 320. Preferably RTSP over TCP (Transmission Control Protocol) or RTSP over another reliable protocol is used in transferring

application layer requests from the client device to the streaming server. The reception of the messages at the streaming server can therefore basically be guaranteed.

- 5 Although it has been described that the cell reselection would be performed between two base stations, it should be noted that a cell reselection can be performed also between two sectors of one and the same base station. Also, it should be noted that depending on the implementation, it may well be that separate messages (PAUSE messages) pausing the sending of streaming media are not needed.
- 10 In an alternative embodiment, both to stop sending streaming media and to start resending streaming media is caused by a single appropriate message.

- Further, in relation to the buffer filling embodiment, it has been described in the foregoing that when the server switches to sending a lower bit rate sequence, the
- 15 original transmission bit rate would be maintained. However, in an alternative embodiment of the invention, a transmission bit rate higher than the original transmission bit rate is used during the filling period in order to fill the buffer more quickly. In this embodiment, it is assumed that a larger bandwidth can be requested by the mobile client device and that a larger bandwidth is actually available in the radio access network.
 - 20

- Further, in relation to the buffer filling embodiment, the bandwidth information of the sequence can be sent to the streaming server 111 also by other means than the RTSP field 'Bandwidth' (for example, by requesting a specific sequence encoded
- 25 at a bitrate known to the client device 101). In this case, the field 'Bandwidth' is not used, the fields 'Speed' and 'Range' are recalculated according to the actual known sequence bit rate.

- It is possible that the network (air-interface) bandwidth changes during the buffer
- 30 filling period. If the client device 101 supports bandwidth adaptation that involves bit stream switching, the client device 101 should pause the buffer filling (with an

RTSP PAUSE message) before sending a bandwidth adaptation message to the streaming server 111. After the bandwidth adaptation operations are over, the client device can start the buffer filling again and recalculate the 'Bandwidth', 'Speed' and 'Range' values according to the bit rate and timing information of the new media stream.

Also, in relation to the buffer filling embodiment, it should be noted that, in an alternative embodiment, the second PAUSE/PLAY messages are not sent at all but the switch back to sending the original media sequence at the original transmission bit rate will be performed automatically by the server based on the stopping point information contained in the first PLAY message.

The resending request mentioned in the various embodiments of the invention may, in certain cases, actually be a sending request. One such case is considered as an alternative embodiment of the invention. In this embodiment, the client device (101) knowing beforehand that a cell reselection is going to occur in a very near future sends a PAUSE message to the streaming server (111) before the beginning of the cell reselection period (i.e. during the pre-CR period). Sending of the PAUSE message is triggered by a cell reselection initiation event notified to the client software 220 by means of the lower layer API. The PAUSE message causes the streaming server 111 to stop sending the streaming media. When the CR-period is over, the client device 101 then sends a PLAY message causing the streaming server to start sending at the point at which sending was stopped before the CR period. Playing of the streaming media is not stopped at the client device 101 in between, and if the client device buffer 240 has been selected to be longer in time than the time that cell reselection takes, the buffer 240 does not become totally empty during cell reselection and the user does not experience any jumps or interruption in playback. The PLAY message may contain a request to send at an increased speed in order to increase the degree of fullness of the buffer 240.

In yet another embodiment, the mobile client device 101 sends, before the

CR period, to the streaming server 111 an appropriate message requesting the streaming server 111 to stop sending the streaming media and to start sending again at a suitable point in time after the CR period. Sending of the message is triggered, during the pre-CR period, by a cell reselection initiation event notified to the client software 220 by means of the lower layer API. The mobile client device 101, knowing approximately the time that the cell reselection is going to take, estimates a suitable point in time at which it is again able to receive data. It inserts this information to the message so as to prevent the streaming server 111 from starting to send again too early.

10

An example of the above mentioned message is as follows:

```
PLAY rtsp://example.com/foo RTSP/1.0
```

```
CSeq: 7
```

15

```
Session: 354832
```

```
Range: npt=28.00-40.00; time=19970123T153600Z
```

```
Bandwidth: 20000
```

```
Speed: 1.5
```

20 This is an RTSP PLAY message having a header field 'Time'. The value of that field schedules the start of the "future" sending of the streaming media.

Figure 11 illustrates (especially relating to the buffer filling embodiments) the following three ways in which messages may be sent:

25

Case 1: The first PAUSE and PLAY messages are sent after the CR period. The second PAUSE and PLAY messages are sent when the buffer 240 is full.

30 Case 2: The first PAUSE message is sent before the CR period. The associated PLAY message is sent after the CR period. The second PAUSE and PLAY messages are sent when the buffer 240 is full.

Case 3: The first PAUSE message is sent before the CR period. The associated PLAY message is sent after the PAUSE message but before the CR period. The second PAUSE and PLAY messages are sent when the buffer 240 is full.

5

The sending of the second PAUSE message is not necessary, if the first PLAY message contains a closed 'Range' field.

Particular implementations and embodiments of the invention have been described. It is clear to a person skilled in the art that the invention is not restricted to details of the embodiments presented above (e.g. message names and message field names), but that it can be implemented in other embodiments using equivalent means without deviating from the characteristics of the invention. The scope of the invention is only restricted by the attached patent claims.

15